

(12) UK Patent Application (19) GB (11) 2 390 919 (13) A

(43) Date of A Publication 21.01.2004

(21) Application No: 0315869.8
(22) Date of Filing: 07.07.2003
(30) Priority Data:
(31) 10195993 (32) 16.07.2002 (33) US

(71) Applicant(s):
Hewlett-Packard Development Company LP
(Incorporated in USA - Texas)
20555 S.H.249, Houston, Texas 77070,
United States of America

(72) Inventor(s):
Zachary Steven Smith
Lee Becker
David Albert Heckman

(74) Agent and/or Address for Service:
Carpmaels & Ransford
43 Bloomsbury Square, LONDON,
WC1A 2RA, United Kingdom

(51) INT CL⁷:
G06F 17/50

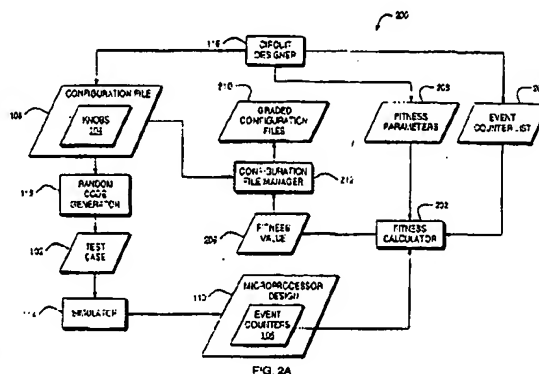
(52) UK CL (Edition W):
G4A AUB

(56) Documents Cited:
US 20020065640 A
AN-6622508; TI-Automatic test bench generation for
validation of RT-level descriptions; an Industrial
experience; PUB-2000; IEEE Comp. Soc, USA 27-30 MAR
2000; AU-Corno F; Sonza Reorda M; Squillero
G; Manzone A; Pincetti A. JSBN
0-7695-0537-6; PG-385-389; ORD-2000.

(58) Field of Search:
UK CL (Edition V) G4A
INT CL⁷ G06F
Other: Online : EPODOC, JAPIO, WPI, FULL TEXT,
INSPEC

(54) Abstract Title: Generating random circuit design tests using a genetic algorithm

(57) Techniques are disclosed for automatically generating test instructions for use in testing a microprocessor design (110). A configuration file (108) includes a plurality of knobs which specify a probability distribution of a plurality of microprocessor instructions. A random code generator (112) takes the configuration file (108) as an input and generates test instructions which are distributed according to the probability distribution specified by the knobs. The test instructions are executed on the microprocessor design (110). The microprocessor behaviors that are exercised by the test instructions are measured and a fitness value (206) is assigned to the configuration file (108) using a fitness function. The configuration file (108) and its fitness value (206) are added to a pool of configuration files (210). A configuration file synthesizer (222) uses a genetic algorithm to synthesize a new configuration file (224) from the pool of existing configuration files (210). This process may be repeated to generate configuration files which increasingly exercise microprocessor behaviors which are of interest.



GB 2 390 919 A

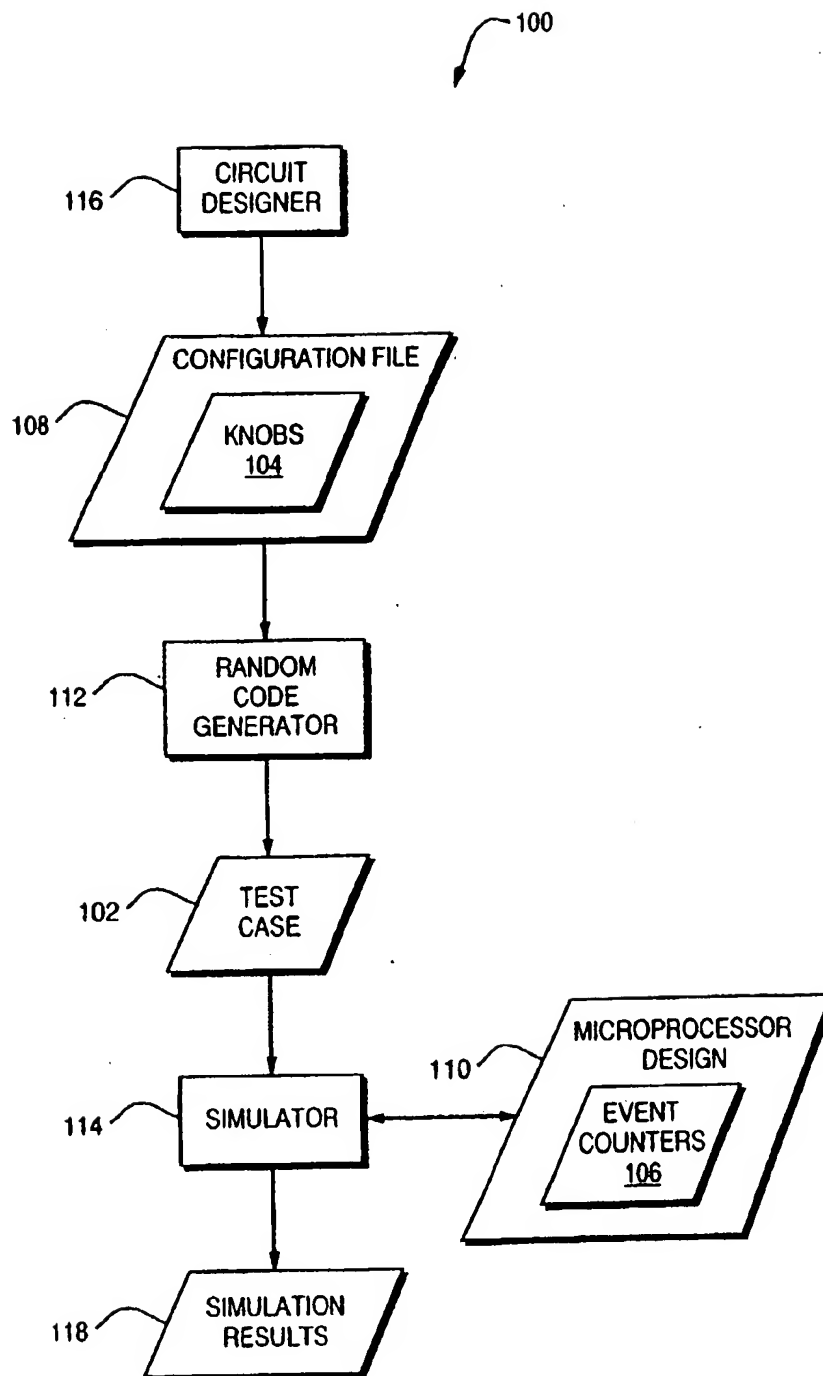


FIG. 1
(PRIOR ART)




FIG. 2A




FIG. 2B


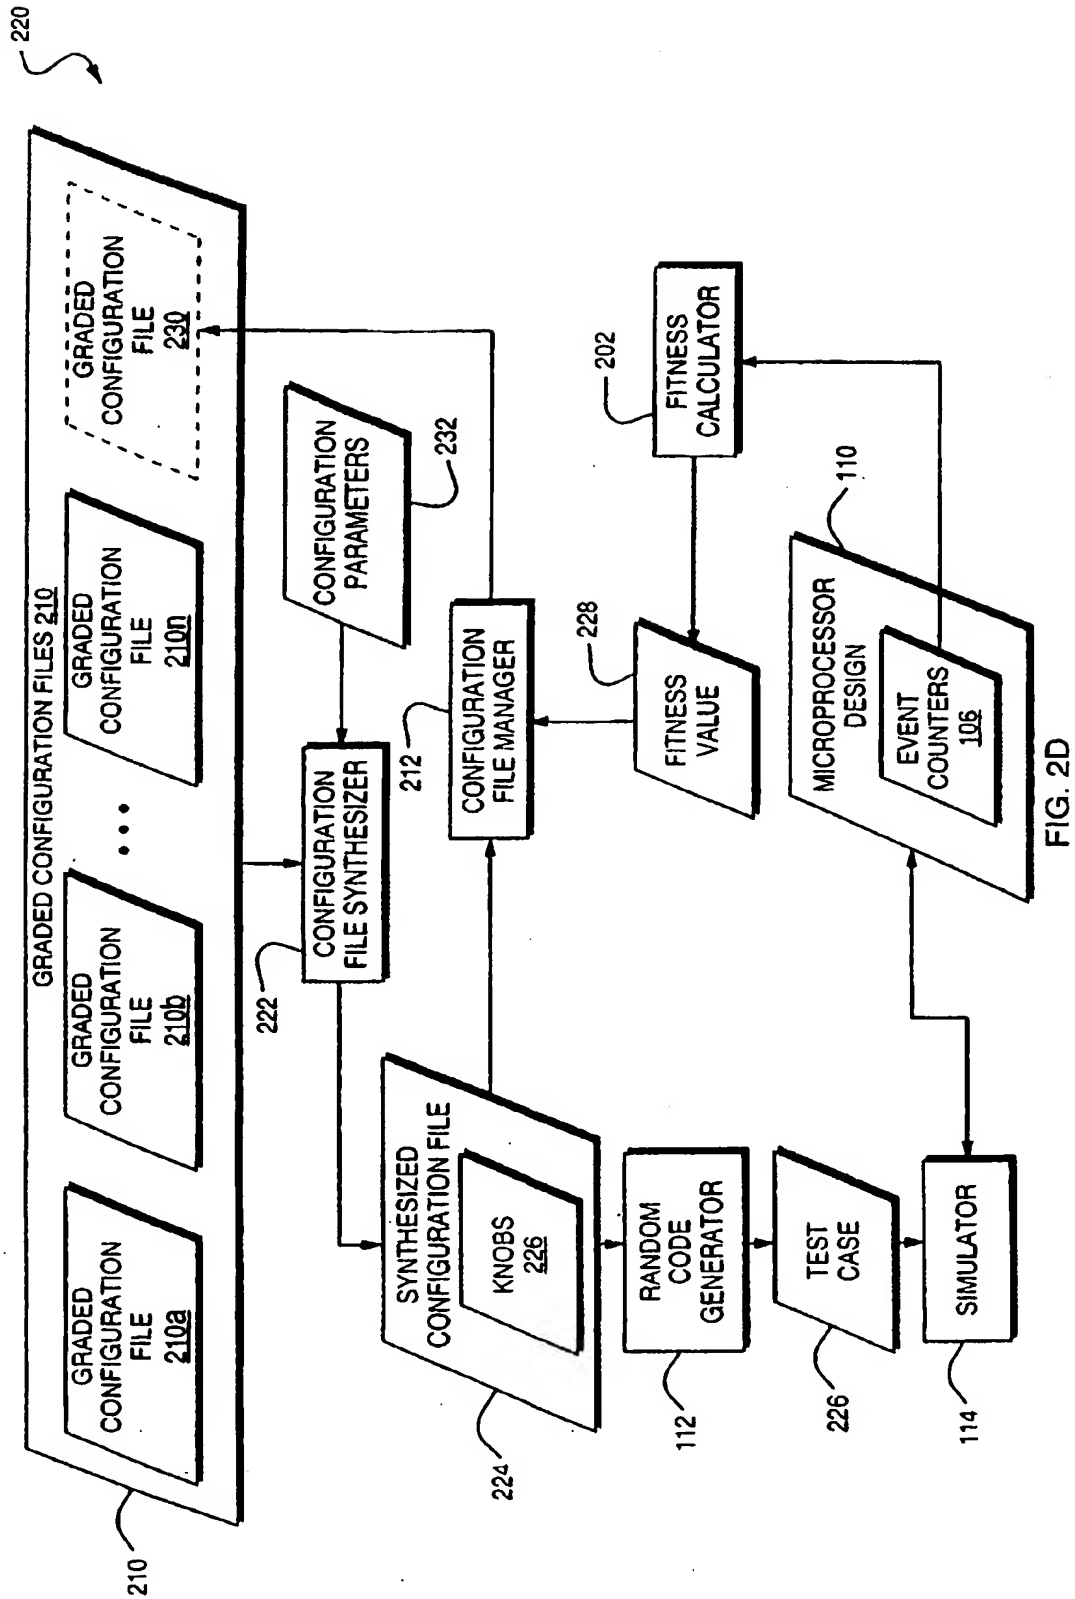


FIG. 2C



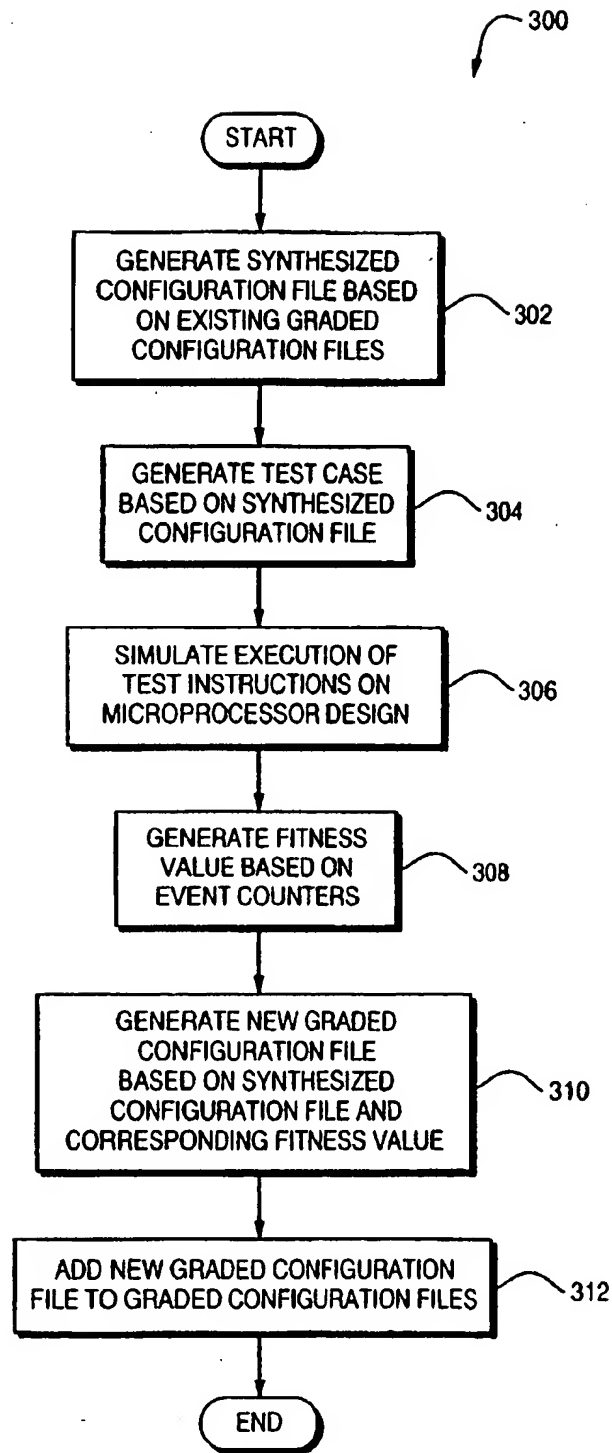


FIG. 3A

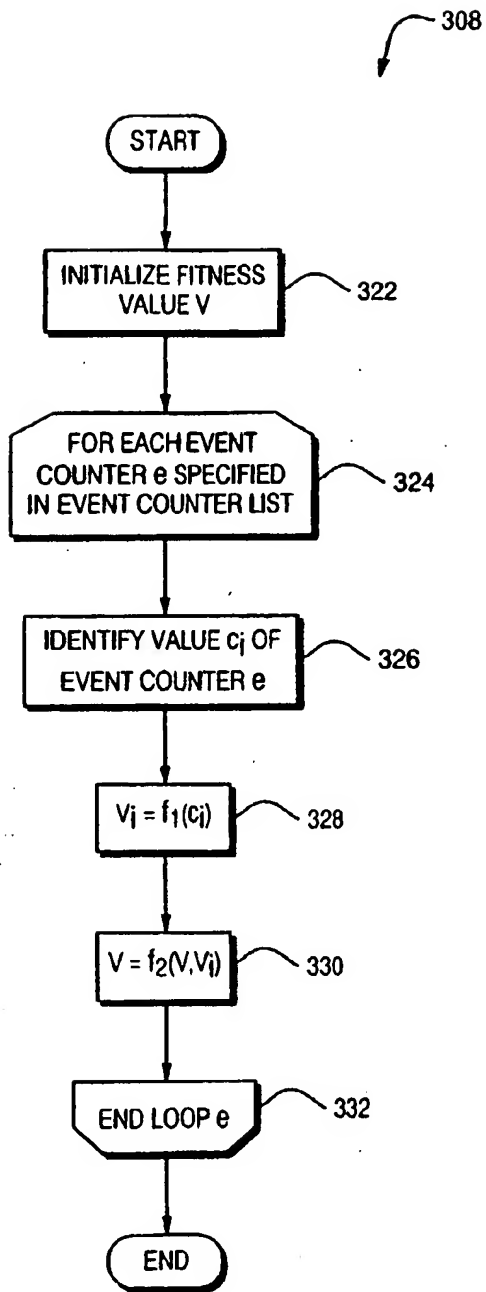


FIG. 3B

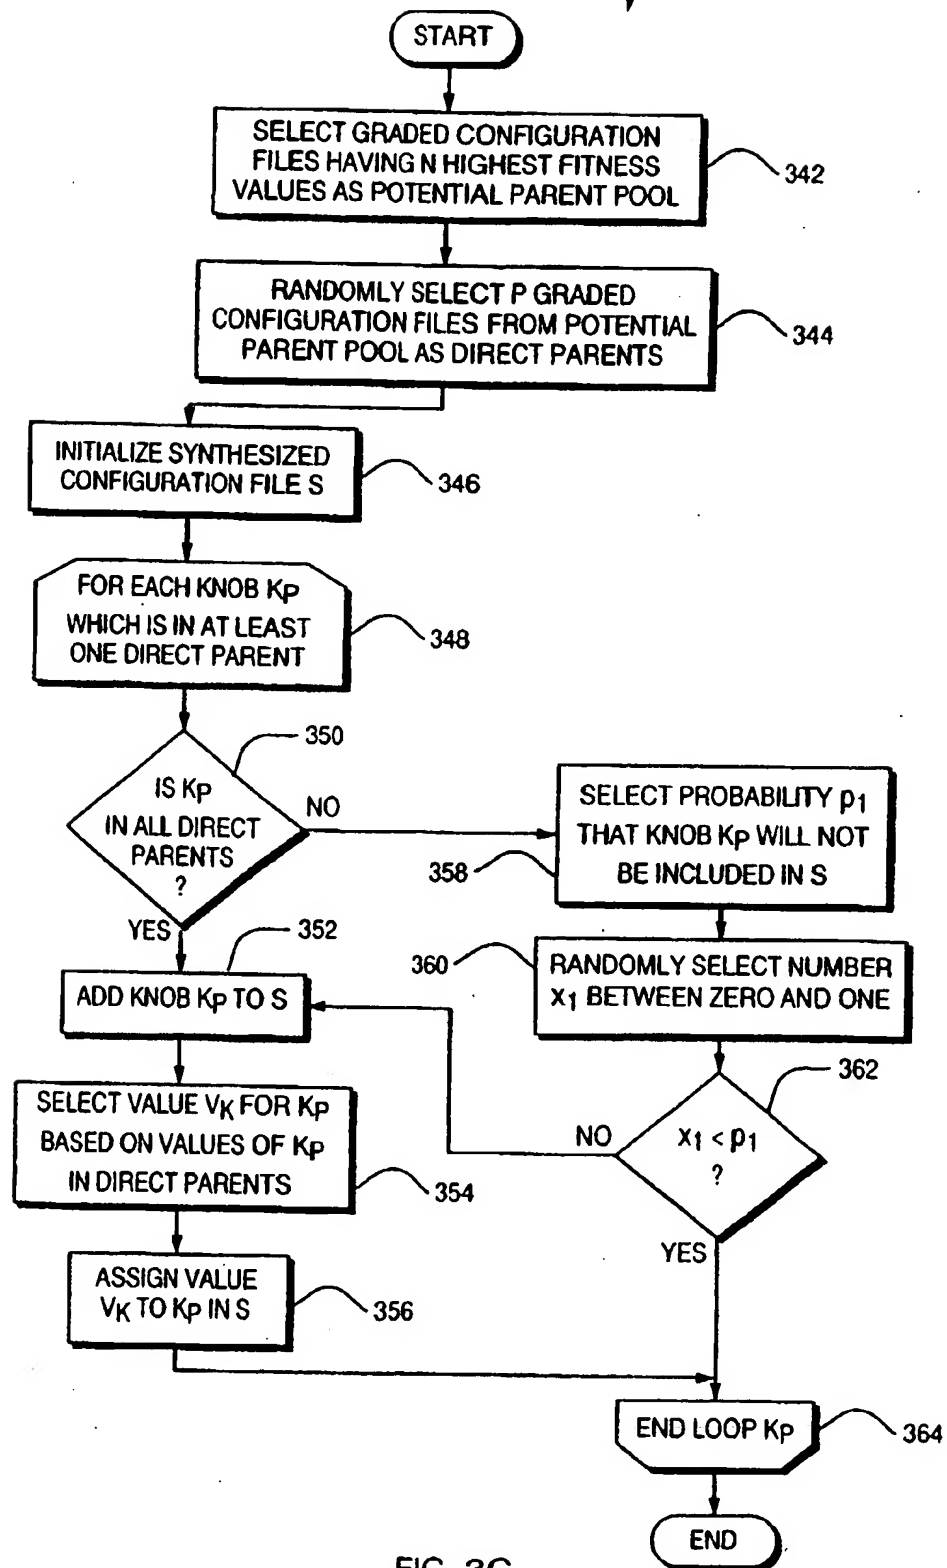


FIG. 3C

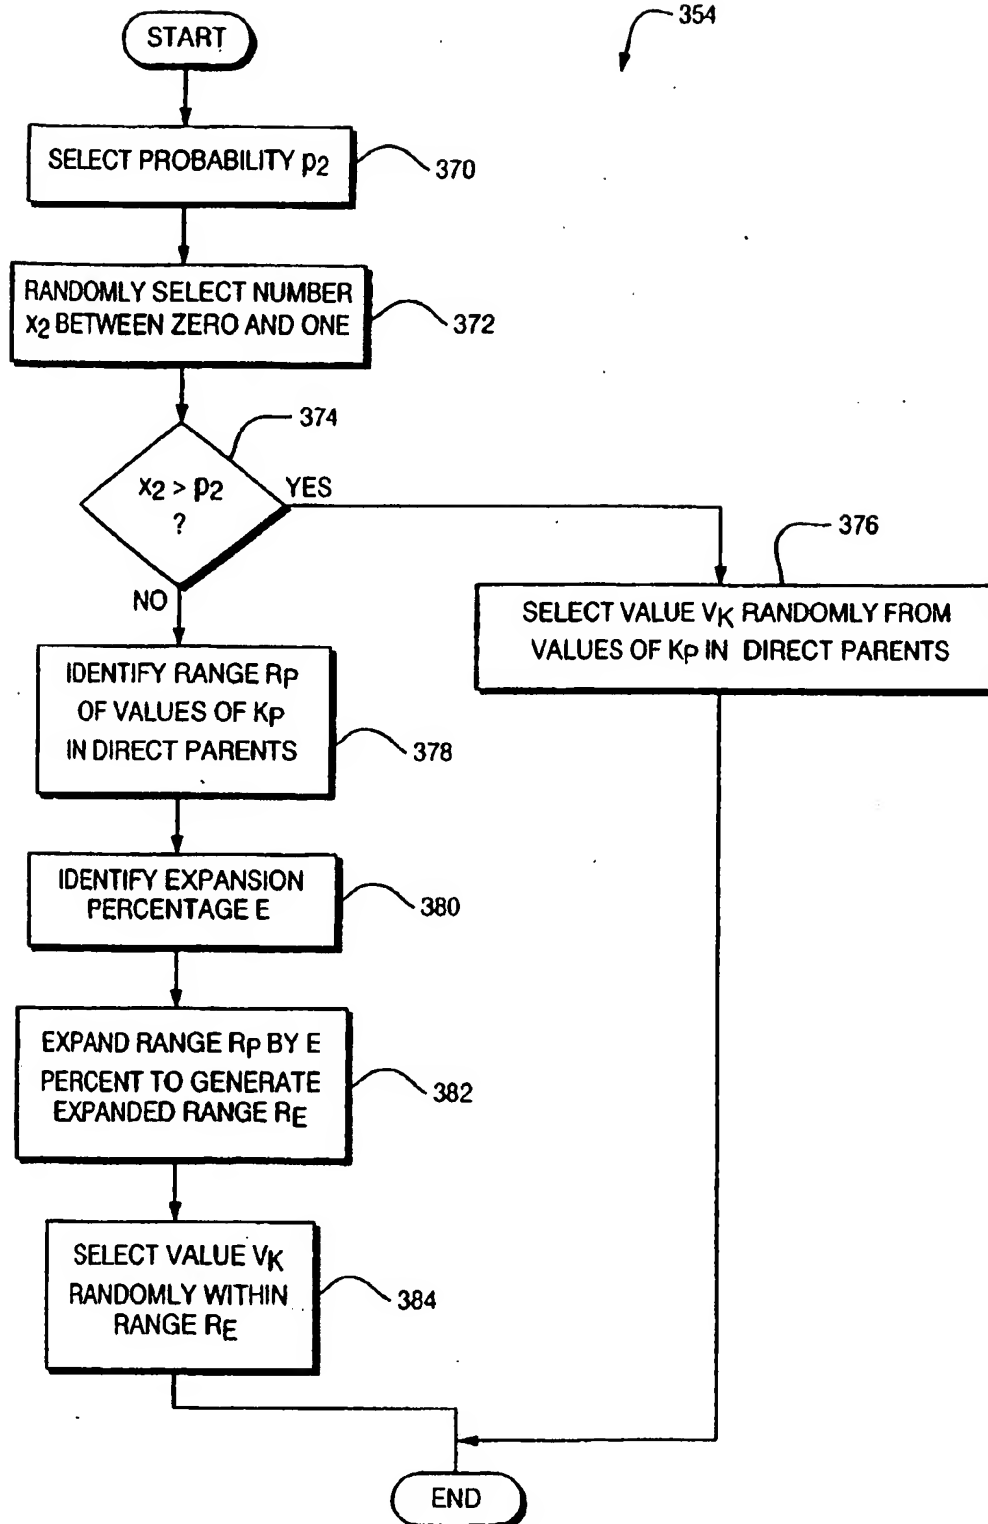


FIG. 3D




FIG. 4




FIG. 5A




FIG. 5B

Random Code Generation Using Genetic Algorithms

BACKGROUND

Field of the Invention

The present invention relates to microprocessor design and, more particularly, to the automatic generation of test instructions for microprocessor designs.

Related Art

Various electronic design automation (EDA) software tools exist for designing microprocessors and other circuitry. Such tools allow circuit designers to create and modify virtual models of the circuit being designed. A circuit designer may, for example, specify a circuit design using a textual description written in a hardware description language (HDL), such as Verilog or VHDL, or by using a graphical user interface to manipulate a graphical representation of the circuit design.

Software tools are also frequently used for testing circuit designs, such as microprocessor designs. Referring to FIG. 1, for example, a prior art microprocessor test system 100 is shown in functional block diagram form. The system includes a microprocessor design 110, also referred to as a hardware model under test or design under test (DUT). Although the microprocessor design 110 may be either an actual (hardware) microprocessor or a software (simulated) model of a microprocessor, assume for purposes of the present discussion that the microprocessor design 110 is a software model of a microprocessor.

The microprocessor design 110 models both the operation of the microprocessor (e.g., the functional relationship between inputs and outputs of the microprocessor) and the

state of the microprocessor's machine resources, such as its registers, cache(s), and main memory, at a particular point in time. The microprocessor design 110 may be implemented, for example, in a data structure in the memory of a computer or in a file stored on a hard disk or other computer-readable medium.

The system 100 also includes a simulator 114, which is typically implemented as a software program. The simulator 114 simulates the operation of the microprocessor modeled by the microprocessor design 110. A significant advantage of using simulators for testing is that they may detect errors prior to the costly and time-consuming fabrication of the microprocessor itself.

The system 100 also includes a test case 102, which specifies both initial values for the (simulated) machine resources of the microprocessor design 110 and test instructions to be executed (in simulation) by the microprocessor design 110. The test case 102 may also specify the outputs that are expected to result from performing the test instructions based on the specified initial values. The test case 102 may, for example, be implemented in a file stored on a computer-readable medium or as a data structure in the memory of a computer.

The simulator 114 receives the test case 102 as an input and initializes the (simulated) machine resources of the microprocessor design 110 with the initial values specified by the test case 102. The simulator 114 then simulates execution of the instructions specified by the test case 102 on the microprocessor design 110. The simulator 114 modifies the state of the microprocessor design 110 accordingly as the simulation progresses. The simulator 114 produces simulation results 118 which indicate, among other things, whether the output produced by executing the test instructions matches the expected output specified by the test case 102.

Although a human circuit designer 116 may create the test case 102 manually, the test case 102 is typically generated automatically by a software program referred to as a random code generator 112. The random code generator 112 creates random sequences of instructions that are intended to exercise the microprocessor design 110 more thoroughly than conventional application software. The random code generator 112 may generate large numbers of test cases automatically and rapidly, thereby facilitating the testing process.

Test cases (such as test case 102) generated by random code generator 112 are typically not, however, completely random. Rather, a good random code generator generates test cases which focus on key aspects of the microprocessor design 110 while retaining enough randomness to test the microprocessor design 110 thoroughly. The circuit designer 116 may use a configuration file 108 (also referred to as a "probability file") to exercise control over which test instructions are generated by the random code generator 112. The configuration file 108 may, for example, specify the frequencies with which different microprocessor instructions are to be simulated by the simulator 114. More specifically, configuration file 108 may include "knobs" 104. Each knob may specify one instruction or class of instructions by name and the frequency with which the instruction or instruction class should occur in the test case 102. The random code generator 112 is designed to generate test cases in which the distribution of instructions and/or instruction classes at least roughly matches the probability distribution specified by the knobs 104.

Microprocessors and microprocessor designs, such as the microprocessor design 110, typically have a performance monitoring unit (PMU) which includes event counters 106 that count the number of times that various events occur in the microprocessor design 110. Examples of events which may be

counted by event counters 106 include, for example, the number of times the microprocessor's current privilege level (CPL) changes, the amount of traffic on the microprocessor's instruction translation lookaside buffer (ITLB) and/or data translation lookaside buffer (DTLB), the number of CPU cycles, the number of retired (executed and committed) null operations, the number of retired load/store operations, and the number of Level 1 and/or Level 2 cache hits and/or misses. The execution of particular test instructions, or the results produced by executing such instructions, may cause one or more events to occur. Each time a particular event occurs, the corresponding one of the event counters 106 is incremented.

When testing the microprocessor design 110 it is often desirable to thoroughly test (or "exercise") the entire range of events. In particular circumstances, however, the circuit designer 116 may wish to focus the simulation on a particular event or events. It is difficult, however, for the circuit designer 116 to design the configuration file 108 to test (or "exercise") the entire range of events or particular events of interest because, for example, the relationship between the execution of particular instruction sequences and the occurrence of particular events may be complex and difficult to predict. Although the circuit designer 116 may improve the extent to which the configuration file 108 exercises particular events by inspecting the event counters 106 and/or the simulation results 118 upon completion of the simulation performed by the simulator 114 and modifying the configuration file 108 manually in response, such a process is tedious, time-consuming, and not guaranteed to improve the extent to which the configuration file 108 exercises events of interest.

What is needed, therefore, are improved techniques for generating test instructions for microprocessor designs.

SUMMARY

Techniques are disclosed for automatically generating test instructions for use in testing a microprocessor design. A configuration file includes a plurality of knobs which specify a probability distribution of a plurality of microprocessor instructions. A random code generator takes the configuration file as an input and generates test instructions based on the probability distribution specified by the knobs. The test instructions are executed on the microprocessor design. The microprocessor behaviors that are exercised by the test instructions are measured and a fitness value is assigned to the configuration file using a fitness function. The configuration file and its fitness value are added to a pool of configuration files. A configuration file synthesizer uses a genetic algorithm to synthesize a new configuration file from the pool of existing configuration files. This process may be repeated to generate configuration files which increasingly exercise microprocessor behaviors which are of interest.

In one aspect, the present invention features a method for use in a system including a circuit design (e.g., a microprocessor design) and a plurality of circuit input vectors (e.g., a plurality of microprocessor test instructions) which have been applied to the circuit design to produce at least one occurrence of at least one event within the circuit design. The method includes steps of: (A) identifying a plurality of counter values indicating a number of occurrences of each of a plurality of events; and (B) applying a fitness function to the plurality of counter values to produce an aggregate fitness value for the plurality of circuit input vectors. The fitness function may, for example, be the sum of the plurality of counter values, or the sum of the expression $1+A*B*\log(c_i)$ for each counter value c_i . The value of A may, for example, be zero and the value of B may,

for example, be one. The circuit design may be a physical circuit or a software model of the circuit. The plurality of events may be events monitored by a performance monitoring unit associated with the circuit design.

In another aspect, the present invention features a method for use in a system including a circuit design and a plurality of configuration files, the plurality of configuration files including a plurality of knobs having a plurality of values specifying a plurality of probability distributions of a plurality of circuit input vectors. The method includes steps of: (A) initializing a synthesized configuration file; (B) selecting a first subset of the plurality of configuration files as a potential parent pool; and (C) for each knob K_p included in at least one of the plurality of configuration files in the potential parent pool, providing a synthesized value for knob K_p in the synthesized configuration file based on at least one value of knob K_p in the plurality of configuration files in the potential parent pool. The plurality of configuration files may be associated with a plurality of fitness values, and the step (B) may include a step of: (B)(1) selecting as the first subset the plurality of configuration files having the highest fitness values in the plurality of configuration files.

The method may further include a step, performed prior to the step (C), of: (D) replacing the plurality of configuration files in the potential parent pool with a subset of the plurality of configuration files in the potential parent pool. The subset may be randomly selected from the plurality of configuration files in the potential parent pool. The step (C) may include a step of: (C)(1) for each knob K_p included in at least one of the plurality of configuration files in the potential parent pool, selecting the synthesized value for knob K_p in the synthesized configuration file from among the at

least one value of knob K_p in the plurality of configuration files in the potential parent pool.

The step (B) may include steps of: (B) (1) for each knob K_p which is included in all of the plurality of configuration files in the potential parent pool, selecting the synthesized value for knob K_p from the at least one value of knob K_p in the plurality of configuration files in the potential parent pool; and (B) (2) for each knob K_p which is not included in all of the plurality of configuration files in the potential parent pool, performing steps of: (B) (2) (a) identifying the ratio of the number of configuration files in the potential parent pool which include the knob K_p to the number of configuration files in the potential parent pool; and (B) (2) (b) determining whether to provide the synthesized value for knob K_p in the synthesized configuration file based on the ratio. The step (B) (2) (b) may include steps of: (B) (2) (b) (i) randomly selecting a number x_1 ; and (B) (2) (b) (ii) deciding to provide the synthesized value for knob K_p in the synthesized configuration file only if x_1 is less than the ratio.

In yet another aspect of the present invention, a method is provided for use in a system including a circuit design and a plurality of configuration files, the plurality of configuration files including a plurality of knobs having a plurality of values specifying a plurality of probability distributions of a plurality of circuit input vectors. The method includes steps of: (A) initializing a synthesized configuration file; (B) selecting a subset of the plurality of configuration files having the highest fitness values as a potential parent pool; (C) randomly selecting a subset of the potential parent pool as a plurality of direct parents; and (D) for each knob K_p included in at least one of the direct parents, providing a synthesized value for knob K_p in the synthesized configuration file based on at least one value of knob K_p in the plurality of direct parents.

The step (D) may include steps of: (D)(1) determining whether knob K_p is included in all of the direct parents; (D)(2) selecting the synthesized value for knob K_p from among the at least one value of knob K_p in the plurality of direct parents if it is determined that knob K_p is included in all of the direct parents; and (D)(3) if it is determined that knob K_p is not included in all of the direct parents, performing steps of: (D)(3)(a) selecting a probability p_1 that the synthesized value for knob K_p will not be included in the synthesized configuration file; (D)(3)(b) randomly selecting a number x_1 ; and (D)(3)(c) including the synthesized value for knob K_p in the synthesized configuration file only if x_1 is less than p_1 . The step (D)(3)(a) may include a step of selecting as the probability p_1 the ratio of the number of direct parents which include the knob K_p to the number of direct parents.

The step (D)(2) may include steps of: (D)(2)(a) selecting a probability p_2 ; (D)(2)(b) randomly selecting a number x_2 ; (D)(2)(c) selecting the synthesized value for knob K_p from among the at least one value of knob K_p in the direct parents if x_2 is greater than p_2 ; and (D)(2)(d) if x_2 is not greater than p_2 , performing steps of: (D)(2)(d)(i) identifying a range R_p of values of knob K_p in the direct parents; and (D)(2)(d)(ii) selecting the synthesized value for knob K_p from within the range R_p . The method may further include a step of: (D)(2)(d)(iii) expanding the range R_p to generate an expanded range R_E ; and the step (D)(2)(d)(ii) may include a step of selecting the synthesized value for knob K_p from within the expanded range R_E .

In another aspect of the present invention, a method is provided for use in a system including a circuit design and a plurality of configuration files, the plurality of configuration files including a plurality of knobs having a plurality of values specifying a plurality of probability distributions of a plurality of circuit input vectors. The

method includes steps of: (A) generating a first synthesized configuration file based on the plurality of configuration files, the first synthesized configuration file specifying a first synthesized probability distribution of a first subset of the plurality of circuit input vectors; (B) using the first synthesized configuration file to generate a first test case comprising the first subset of the plurality of circuit input vectors; (C) providing the first subset of the plurality of circuit input vectors as inputs to the circuit design; (D) generating a fitness value based on at least one event counter in the microprocessor design, the at least one event counter comprising a count of at least one occurrence of at least one event in the microprocessor design in response to provision of the first subset of the plurality of circuit input vectors; (E) generating a graded configuration file including the synthesized configuration file and the fitness value; and (F) adding the graded configuration file to the plurality of configuration files.

The method may further include steps of: (G) after the step (F), generating a second synthesized configuration file based on the plurality of configuration files, the second synthesized configuration file specifying a second synthesized probability distribution of a second subset of the plurality of circuit input vectors; (H) using the second synthesized configuration file to generate a second test case comprising the second subset of the plurality of circuit input vectors; and (I) providing the second subset of the plurality of circuit input vectors as inputs to the circuit design.

Other features and advantages of various aspects and embodiments of the present invention will become apparent from the following description and from the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a data flow diagram of a prior art system for testing a microprocessor design;

FIG. 2A is a data flow diagram of a system for testing a microprocessor design according to one embodiment of the present invention;

FIG. 2B is a diagram of a stream of configuration files and corresponding fitness values generated by a fitness calculator according to one embodiment of the present invention;

FIG. 2C is a diagram of a plurality of graded configuration files according to one embodiment of the present invention;

FIG. 2D is a data flow diagram of a system for synthesizing a new configuration file from existing graded configuration files according to one embodiment of the present invention;

FIG. 3A is a flowchart of a method performed by the system of FIG. 2D for synthesizing a new configuration file from existing graded configuration files, assigning a grade to the synthesized configuration file, and adding the synthesized configuration file and corresponding grade to the graded configuration files according to one embodiment of the present invention;

FIG. 3B is a flowchart of a method for calculating fitness value for a configuration file according to one embodiment of the present invention;

FIG. 3C is a flowchart of a method for synthesizing a new configuration file from existing configuration files according to one embodiment of the present invention;

FIG. 3D is a flowchart of a method for selecting a value for a knob in a synthesized configuration file according to one embodiment of the present invention;

FIG. 4 is a data flow diagram of a system for generating a fitness value for a configuration file according to one embodiment of the present invention; and

FIGS. 5A-5B are graphs of fitness values generated by embodiments of the present invention for sequences of configuration files.

DETAILED DESCRIPTION

Referring to FIG. 2A, a system 200 for testing the microprocessor design 110 according to one embodiment of the present invention is shown. The system 200 includes configuration file 108, random code generator 112, test case 102, simulator 114, and microprocessor design 110, which may, for example, be conventional elements as described above with respect to FIG. 1. The system 200 also includes however, a fitness calculator 202 which may, for example, be implemented as a software program. Upon completion of the simulation performed by the simulator 114, the fitness calculator 202 may generate a fitness value 206 (also referred to herein as a "grade") which indicates the extent to which the test instructions in the test case 102 exercise events or other behaviors in the microprocessor design 110 that are of interest to the circuit designer 116. The fitness value 206 therefore indirectly indicates the fitness of the knobs 104 in the configuration file 108.

The circuit designer 116 may, for example, provide an event counter list 204 to the fitness calculator 202. The event counter list 204 may specify one or more of the event counters 106 which count events that are of interest to the circuit designer 116. The fitness calculator 202 generates the fitness value 206 based on the values of the event counters specified in the event counter list 204. The event counter list 204 may identify particular ones of the event counters in any manner. For example, each of the event

counters 106 may have a unique identifier, such as a numerical or textual identifier, in which case the event counter list 204 may identify event counters using their unique identifiers.

The circuit designer 116 may also provide additional fitness parameters 208 to the fitness calculator 202 for use in generating the fitness value 206. Particular examples of the fitness parameters 208 are described below with respect to FIG. 4. Particular examples of techniques that the fitness calculator 202 may use to generate the fitness value 206 are described in more detail below with respect to FIG. 3B.

A configuration file manager 212 adds the configuration file 108, along with its associated fitness value 206, to a collection of graded configuration files 210. If the circuit designer 116 provides a new configuration file to the random code generator 112, the fitness calculator 202 will generate a new fitness value corresponding to the new configuration file. The configuration file manager 212 will add the new configuration file, along with its associated fitness value 206, to the graded configuration files 210.

For example, referring to FIG. 2B, a stream of configuration files 108a-n is shown along with corresponding fitness values 206a-n, in which n is the number of configuration files and corresponding fitness values. Each of the configuration files 108a-n is used by the random code generator 112 to generate corresponding test cases. Although FIG. 2A shows only one test case 102 corresponding to the configuration file 108, the random code generator 112 typically generates many test cases from a single configuration file. Returning to FIG. 2B, for each of the configuration files 108a-n: (1) the simulator 114 simulates execution of each of the corresponding test cases on the microprocessor design 110, and (2) the fitness calculator 202 generates a corresponding one of the fitness values 206a-n.

In other words, upon execution of all of the test cases corresponding to a particular one of the configuration files 108a-n, the fitness calculator 202 generates a fitness value for the configuration file. Although the configuration files 108a-n are illustrated as inputs to the fitness calculator 202 in FIG. 2B for ease of illustration, in practice the information contained in the configuration files 108a-n is only input to the fitness calculator 202 indirectly through the sequence of events illustrated in FIG. 2A.

Referring to FIG. 2C, an example of the graded configuration files 210 is shown. In the example illustrated in FIG. 2C, graded configuration files 210 include graded configuration files 210a-n, each of which includes one of the configuration files 108a-n and a corresponding one of the fitness values 206a-n. For example, graded configuration file 210a includes configuration file 108a and the corresponding fitness value 206a. Although in the example described above the configuration files 108a-n are generated manually by the circuit designer 116, one or more of the configuration files 108a-n may be generated automatically, as described in more detail below.

Referring to FIG. 2D, a system 220 for synthesizing a new configuration file 224 from existing graded configuration files 210a-n is shown according to one embodiment of the present invention. Referring to FIG. 3A, a flowchart is shown of a method 300 that is performed by the system 220 according to one embodiment of the present invention. A configuration file synthesizer 222 takes as input the existing graded configuration files 210a-n (illustrated in solid outline in FIG. 2D) and synthesizes from them the synthesized configuration file 224, containing knobs 226 (step 302). In particular, the configuration file synthesizer 222 may synthesize the synthesized configuration file 224 from the existing graded configuration files 210a-n having the highest

fitness values. As described in more detail below, the existing graded configuration files having the highest fitness values are referred to as the "potential parent pool." Various techniques that the configuration file synthesizer 222 may use to generate the synthesized configuration file 224 will be described in more detail below with respect to FIG. 3C-5D.

The synthesized configuration file 224 is provided as input to the random code generator 112, which generates test case 226 based on the synthesized configuration file 224 in the manner described above with respect to FIG. 2A (step 304). Although only one test case 226 is shown in FIG. 2D, in practice the random code generator 112 may generate many test cases based on the synthesized configuration file 224. The simulator 114 simulates execution of the instructions in the test case 226 (and any other test cases which have been generated) on the microprocessor design 110, thereby causing the values of event counters 106 to change appropriately (step 306). The fitness calculator 202 generates a fitness value 228 for the synthesized configuration file 224 (step 308). (The fitness parameters 208 and the event counter list 204, shown in FIG. 2A, are also input to the fitness calculator 202 but are omitted from FIG. 2D for ease of illustration.)

The configuration file manager 212 generates a graded configuration file 230 corresponding to the synthesized configuration file 224 and its corresponding fitness value 228 (step 310), and adds the graded configuration file 230 to the graded configuration files 210 (step 312). In this way, the synthesized configuration file 224 becomes available for use in synthesizing additional synthesized configuration files (by repeatedly performing method 300), which may in turn be used to synthesized additional synthesized configuration files, and so on.

As described in more detail below, the configuration file synthesizer 222 may synthesize the most fit ones of the existing graded configuration files 210a-n in such a manner that the knobs 226 in the synthesized configuration file 224 specify an instruction probability distribution that results in test cases which tend to better exercise the events that are of interest to the circuit designer 116 than the existing graded configuration files 210a-n. By generating synthesized configuration files in this way and using them as inputs to the random code generator 112, the system 220 tends to increasingly exercise the events that are of interest to the circuit designer 116.

Having described in general overview the operation of various embodiments of the present invention, embodiments of the present invention will now be described in more detail. For example, referring to FIG. 3B, a flow chart is shown of a method 308 that is performed by the fitness calculator 202 to calculate the fitness value 206 (referred to below by the variable V) according to one embodiment of the present invention.

The method 308 initializes the fitness value V (step 322). The fitness value V may be initialized to any appropriate value, such as zero or one. The method 308 enters a loop over each event counter e specified in the event counter list 204 (step 324). As described above, the event counter list 204 may specify a subset of the event counters 106 that is of interest to the circuit designer 116. The method 308 identifies the value c_i of event counter e (step 326). As described above, the value c_i represents the number of times that the corresponding event occurred while executing the test instructions in the test case 102.

The fitness calculator 202 may include a fitness function. For example, referring to FIG. 4, one embodiment of the fitness calculator 202 is illustrated in more detail. As

shown in FIG. 4, the fitness calculator includes a fitness function 408. In the particular embodiment illustrated in FIG. 4, the fitness parameters 208 include two parameters 402 and 404, designated by variables A and B , respectively. Specified event counters 406 include the values of the event counters specified in the event counter list 204. Specified event counters 406 includes event counter values $c_0, c_1, c_2, \dots, c_m$, where m is the number of event counters specified in the event counter list 204. The method 308 may, for example, generate the list of specified event counters 406 prior to execution of the method 308 or retrieve each of the specified event counter values from the event counters 106 as needed.

The fitness function 408 may be any function. For example, the fitness function may simply be the sum of each of the counter values c_i . Such a fitness function provides a straightforward measure of fitness by making the fitness value equal to the total number of times that any event specified in the event counter list occurs during simulation.

Referring to FIG. 4, the fitness function alternatively may be the sum of the expression $(1+A*B*\log(c_i))$ for each counter value c_i . Such a function reflects a relatively strong interest in any event counter which has a non-zero value (corresponding to events which occurred at least once during simulation) and a relatively weak interest in additional occurrences of such events. In the event that the circuit designer 116 does not specify values for one or more of the fitness parameters 208, the fitness calculator 202 may provide default values. For example, in one embodiment default value for the A parameter is zero and the default value for the B parameter is one.

The fitness function 408 may, for example, include both: (1) a counter fitness function f_1 which generates a counter fitness value V_i based on the value of the current event counter c_i , and (2) a combination function f_2 for combining the

results of the counter fitness function f_i for each of the error counter values. For example, returning to FIG. 3B, the method 308 may calculate counter fitness value V_i based on the counter value c_i using counter fitness function f_i (step 328). The method 308 may update the value of V based on the current value of V and the value of V_i using the combination function f_2 (step 330). The method 308 repeats steps 326-330 (step 332) to generate the final fitness value V , which is output as fitness value 206.

For example, in the embodiment illustrated in FIG. 4, the counter fitness function $f_i(c_i) = 1 + A + B \cdot \log(c_i)$ and the combination function $f_2(V, V_i) = V + V_i$. In other words, if the fitness function 408 generates the fitness value 206 by summing the values of the specified event counters 406, the counter fitness function $f_i(c_i) = c_i$ and the combination function $f_2(V, V_i) = V \cdot V_i$.

The method 308 illustrated in FIG. 3B may be applied to generate the fitness value 206 when the test case 102 is either a single test case or a plurality of test cases. If the test case 102 is a single test case, the method 308 may be applied to generate the fitness value 206 based on the values of the event counters 106 after execution of the test case 102. If the test case 102 is a plurality of test cases, the values of the event counters 106 after execution of each test case may be combined in any of a variety of ways prior to being input to the method 308. For example, in one embodiment, the method 308 may be performed after the execution of each test case to generate a fitness value for each test case. The fitness values for each test case may then be combined by addition or some other function to generate the overall fitness value 206 for the configuration file 108.

As described above with respect to FIG. 3A, the configuration file synthesizer 222 may generate the

synthesized configuration file 224 based on the existing graded configuration files 210a-n (FIG. 3A, step 302). Referring to FIG. 3C, a flow chart is shown of a method that may be used to implement step 302. The method 302 selects the N graded configuration files having the highest fitness values (step 342). The configuration files selected in step 342 are referred to herein as the "potential parent pool." The configuration file synthesizer 222 may perform step 342 by, for example, sorting the graded configuration files 210a-n in descending order of fitness value and selecting the first N graded configuration files in the sorted list as the potential parent pool. The configuration file synthesizer 222 may select the value N in any manner. Referring to FIG. 2D, for example, the circuit designer 116 may provide the value of N to the configuration file synthesizer 222 as one of a plurality of configuration parameters 232.

Returning to FIG. 3C, the method 302 randomly selects P graded configuration files from the potential parent pool (step 344). The graded configuration files selected in step 344 are referred to herein as "direct parents." The configuration file synthesizer 222 may select the value P in any manner. For example, the circuit designer 116 may provide the value of P as one of the configuration parameters 232. As described in more detail below, the configuration file synthesizer 222 synthesizes knob values from the direct parents to generate knob values for the synthesized configuration file 224.

As described above, the knobs 104 may specify individual instructions or instruction classes and corresponding probabilities. Each knob typically consists of a knob name and corresponding knob value. The knobs in a configuration file may therefore be represented as a list of knob name-value pairs. A knob name may, for example, be the name or other identifier of an instruction or a class of instructions.

Examples of instruction classes include, but are not limited to, data transfer instructions, Arithmetic-Logic Unit (ALU) instructions, string instructions, floating point instructions, flag manipulation instructions, system instructions, semaphore instructions, Intel MMX™ instructions, instruction set switching branches, inter-segment branches, and intra-segment branches.

Returning to FIG. 3C, the method 302 initializes the synthesized configuration file 224 (step 346). Initializing the synthesized configuration file 224 may include, for example, creating an empty list of knob name-value pairs or creating a list of knob name-value pairs in which the values are undefined (such as null values). According to the description herein, a configuration file "includes" a knob if the configuration file includes a defined value for the knob. If, for example, a configuration file includes a knob name-value pair for a particular knob, but the value of the name-value pair is undefined (such as a null value), then the configuration file does not "include" the knob as that term is used herein. Similarly, a configuration file does not "include" a knob if the configuration file does not include a name-value pair for the knob.

The method 302 enters a loop over each knob K_p which is included in at least one of the direct parents (step 348). The method 302 determines whether the knob K_p is in all of the direct parents (step 350). If it is, the method 302 adds knob K_p to the synthesized configuration file 224 (step 352). As described above, each knob in a configuration file typically consists of a knob name and corresponding knob value. The knobs 226 in the synthesized configuration file 224 may therefore be represented as a list of knob name-value pairs. The method 302 may "add" knob K_p to the synthesized configuration file 224 in step 352 by, for example, adding a

name-value pair having the name of knob K_p to the knobs 226 in the synthesized configuration file 224.

The method 302 selects a value V_K for knob K_p based on the values of knob K_p in the direct parents (step 354). The method 302 may select the value V_K in any of a variety of ways, one example of which is described below with respect to FIG. 3D. The method 302 assigns the value V_K to the knob K_p in the knobs 226 of the synthesized configuration file 224 (step 356). The method 302 may, for example, perform step 356 by assigning value V_K to the value of the name-value pair for knob K_p in the knobs 226.

If the method 302 determines that the knob K_p is not in all direct parents in step 350, then the knob K_p is in at least one but not all of the direct parents. When a configuration file does not include (e.g., does not specify a value for) a particular knob, the random code generator 112 may provide a default value for that knob. Typically it is not possible, however, to determine ahead of time which default value the random code generator 112 will assign to a particular knob. In one embodiment, if a particular knob is included in at least one but not all of the direct parents, the method 302 excludes the knob from the synthesized configuration file 224 with a probability that is directly proportional to the percentage of direct parents which do not include the knob.

For example, referring again to FIG. 3C, if the knob K_p is in at least one (step 348) but not all (step 350) of the direct parents, the method 302 selects a probability p_i that knob K_p will not be included in the knobs 226 of the synthesized configuration file 224 (step 358). The configuration file synthesizer 222 may select the value p_i in any manner. For example, the method 302 may calculate the probability p_i as the percentage of direct parents which do not contain the knob K_p (i.e., the number of direct parents which do not contain the knob K_p divided by the number P of direct

parents). Alternatively, referring to FIG. 2D, the circuit designer 116 may, for example, include the value of p_i as one of the configuration parameters 232. Assuming for purposes of example that the probability p_i is expressed as a decimal number between zero and one, the method 302 randomly selects a number x_i between zero and one (step 360). If $x_i < p_i$, (step 362), the method 302 adds knob K_p to the knobs 226 of the synthesized configuration file 224 (steps 352-356) as described above. Otherwise, the method 302 does not add knob K_p to the knobs 226 of the synthesized configuration file 224. The effect of steps 358-362 is that, in the case of any knob which is in at least one but fewer than all of the direct parents, there is a probability p_i that the knob will not be added to the knobs 226 of the synthesized configuration file 224. The purpose of steps 358-362 is to make the probability that a particular knob will be included in the synthesized configuration file 224 the same as the probability that the knob will be found in a particular one of the direct parents.

The method 302 repeats steps 350-362 for the remaining knobs K_p which are in at least one of the direct parents (step 364). Upon the completion of method 302, the knobs 226 include knobs whose values have been synthesized from the values of knobs in the most fit ones of the graded configuration files 210a-n. As described above with respect to FIG. 2D, the synthesized configuration file 224 generated by the method 302 may be used to test the microprocessor design 110 and to generate yet another graded configuration file which may be added to the graded configuration files 210. In this way, each new synthesized configuration file generated by the configuration file synthesizer 222 is likely to have a higher fitness value than previous configuration files and therefore to better exercise the features of the microprocessor design 110 that are of interest to the circuit designer 116 than the previous configuration files.

As described above with respect to FIG. 3C, the configuration file synthesizer 222 may select a value V_K for the knob K_P based on the values of knob K_P in the direct parents (step 354). Referring to FIG. 3D, a flow chart is shown of a method 354 that is used to select the value V_K according to one embodiment of the present invention.

The method 354 selects a probability p_2 (step 370). The method 354 may select the value p_2 in any manner. Referring to FIG. 2D, for example, the circuit designer 116 may provide the value of p_2 as one of the configuration parameters 232.

Assuming for purposes of example that the probability p_2 is a decimal number between zero and one, the method 354 randomly selects a number x_2 between zero and one (step 372). The method 354 determines whether $x_2 > p_2$ (step 374). If it is, the method 354 selects the value V_K randomly from the values of knob K_P in the direct parents (step 376).

If x_2 is not greater than p_2 , the method 354 identifies the range R_P of values of knob K_P in the direct parents (i.e., the difference between the maximum and minimum values of knob K_P in the direct parents) (step 378). The method 354 identifies an expansion percentage E (step 380). The method 354 may select the value E in any manner. Referring to FIG. 2D, for example, the circuit designer 116 may provide the value of E as one of the configuration parameters 232. The method 354 expands the range R_P by E percent to generate an expanded range R_E (step 382). The method 354 may, for example, expand the range R_P at the bottom of range R_P , at the top of range R_P , or a combination of both to generate the expanded range R_E . The method 354 selects the value of V_K randomly within the expanded range R_E (step 384).

Selection of the value V_K is complete upon completion of step 376 or 384. Once the value V_K is selected, control returns to step 356 of the method 302 illustrated in FIG. 3C.

According to the method 354, there is a probability of p_2 that the value of knob K_p will be chosen at random in the range R_k , and a probability of $1-p_2$ that the value of knob K_p will be chosen randomly from the values of knob K_p in the direct parents. Advantages of choosing the knob value in this way will be described below.

As described above, the fitness values 206a-n (FIG. 2C) that are assigned to the configuration files 108a-n reflect the degree to which the test instructions generated based on the configuration files 108a-n exercise the events that are of interest to the circuit designer 116. If the techniques disclosed herein are implemented appropriately, the fitness values of new synthesized configuration files generated by the configuration file synthesizer 222 should tend to increase over time.

For example, referring to FIGS. 5A-5B, graphs are shown of the fitness values of a sequence of configuration files CF_i over time, in which the index i is referred to herein as a "configuration file number." For example, the first configuration file in the sequence has configuration file number zero, the second configuration file number has configuration file number one, and so on. In both FIGS. 5A and 5B, the x axis represents the configuration file number, and the y axis represents the fitness value of the corresponding configuration file. In the particular examples illustrated in FIG. 5A-5B, the fitness function illustrated in FIG. 4 was used to calculate fitness values, in which $A = 0$ and $B = 1$. The graphs shown in FIGS. 5A-5B are merely examples, and the actual results obtained in particular implementations may vary depending upon the details of the implementation, such as the particular configuration parameters 232, fitness parameters 208, and event counter list 204 that are used.

Referring to FIG. 5A, a graph is shown in which all configuration files in the sequence were generated using the configuration file synthesizer 222. As may be seen from FIG. 5A, the fitness values of the configuration files generated by the configuration file synthesizer 222 generally increase over time, leveling off after the generation of approximately eighty configuration files.

Referring to FIG. 5B, a graph is shown in which approximately the first 480 configuration files (i.e., from $x=0$ through $x=480$ in the graph) were generated manually using conventional techniques, and in which the configuration file synthesizer 222 was used to generate configuration files beginning at approximately $x=480$, using the previous 480 configuration files as a starting point. As may be seen from FIG. 5B, the fitness values of the configuration files increased rapidly upon commencing use of the configuration file synthesizer 222, with fitness values leveling off at approximately $x=700$.

Among the advantages of the invention are one or more of the following.

It is difficult to manually design configuration files which exercise behaviors (e.g., events) in the microprocessor design 110 which are of interest to the circuit designer 116. One reason for this difficulty is that the relationship between the test instruction probabilities that are specified in configuration files and the behaviors which are exercised by the resulting test instructions is complex and difficult to predict. For example, a particular event may occur as the result of a long and/or complex sequence of instructions, and the relationship between the event and the sequence of instructions may not be apparent. The techniques described above in effect both: (1) identify the relationship between test instructions and events, and (2) create new configuration files which more closely specify test instruction

distributions which will exercise the events of interest to the circuit designer 116. The techniques disclosed herein use the actual event counter values resulting from the application of a particular configuration file as feedback which is used to generate subsequent configuration files. This feedback is applied automatically, rapidly, and repeatedly to improve the ability of configuration files to exercise the behaviors that are of interest to the circuit designer 116, thereby relieving the circuit designer 116 of the responsibility for manually inspecting the event counters 106 and determining how to modify the configuration file in response.

Although the use of a conventional random code generator, by itself, does represent an improvement over purely manual generation of code, conventional random code generators do not provide any feedback about the effectiveness of the code they produce for testing events of interest. Even when a conventional random code generator is used, it is typically necessary for the circuit designer 116 to manually inspect the simulation results 118 and/or the event counters 106 manually and attempt to modify the configuration file 108 in response. This can be a tedious and time-consuming process.

The techniques disclosed herein subject the configuration files 210 to an evolutionary process which identifies the configuration files which are most "fit" and which selects knob values from the fittest configuration files for use in subsequent configuration files. This is an example of what is referred to as a "genetic algorithm." Furthermore, the techniques disclosed herein are not limited to using knob values which explicitly appear in the direct parent configuration files. Rather, as described above with respect to FIG. 3D, in some cases the configuration file synthesizer 222 "mutates" knob values from the direct parents to produce a knob value in the new graded configuration file 230 which does not exist in any of the direct parents (FIG. 3D, step 378-

384). This use of mutation strikes a balance between breadth of behavior coverage and focusing on desired behaviors.

By improving the fitness of the configuration files that are used to test the microprocessor design 110 over time, the techniques disclosed herein increase the likelihood that interesting behaviors and decreases the likelihood that uninteresting behaviors will be tested. The techniques disclosed herein in effect guide the random code generator 112 to generate instructions that exercise certain behaviors that are of interest to the circuit designer 116.

It is to be understood that although the invention has been described above in terms of particular embodiments, the foregoing embodiments are provided as illustrative only, and do not limit or define the scope of the invention. Various other embodiments, including but not limited to the following, are also within the scope of the claims.

Although the techniques disclosed herein as described with respect to the microprocessor design 110, the same techniques may be applied to other kinds of circuit designs. More generally, for example, the test case 102 may be any set of circuit input vectors provided to a circuit design. Similarly, the event counters 106 may be any output produced by the circuit design. Alternatively, the event counters 106 may, for example, be internal signals or resource values (such as the values of memory locations) within a circuit design. The fitness calculator 202 may generate the fitness value 206 based on any function of the event counters.

Furthermore, although the examples disclosed herein describe the microprocessor design 110 as a software model of a microprocessor, the techniques disclosed herein may also be applied to physical microprocessors. The event counters 106 may, for example, be hardware rather than software event counters, and the test instructions in the test case 102 may

be provided directly to the microprocessor design 110 for execution without the use of the simulator 114.

The various knobs described herein (such as the knobs 104 and the knobs 226) may each correspond to a particular instruction or to a class (type) of instruction. Similarly, each of the events counted by the event counters 106 may either be an individual event or a class (type) of event. More generally, the term "event" refers herein generally to any behavior of the microprocessor design 110 which may be characterized by a numerical value or other data structure. The event counters 106 are not limited, for example, to counting events which are typically counted by hardware PMUs. Furthermore, the event counters 106 may count occurrences of event which do not strictly occur within the microprocessor design 110 but which are derived from events which occur within the microprocessor design 110. For example, the event counters 106 may count events generated by external analysis tools which analyze the operation of the microprocessor design 110.

Although the drawings illustrate various data structures (e.g., the configuration file 108, test case 102, fitness value 206, fitness parameters 208, event counter list 208, graded configuration files 210, and configuration parameters 208) as having particular logical structures, these are provided merely for purposes of example and do not constitute limitations of the present invention. Rather, alternative data structures for representing equivalent information and for performing equivalent functions will be apparent to those of ordinary skill in the art. Furthermore, although various data structures are described as being implementable as files in a computer file system, this is not a limitation of the present invention. Rather, such data structures may be implemented as binary files, database files, tables generated by word processing or spreadsheet software, header files or

other source code files written in a programming language such as C, or using any appropriate data structure stored on any appropriate computer-readable medium. In particular, the term configuration "file" refers not only to a file stored in a computer file system, but more generally to any computer-implemented means for specifying a probability distribution of test instructions to execute on the microprocessor design 110.

Elements and components described herein may be further divided into additional components or joined together to form fewer components for performing the same functions. For example, the fitness calculator 202, configuration file manager 212, random code generator 112, simulator 114, and configuration file synthesizer 222 may be combined and/or separated into one or more components.

Although the examples described herein are described in relation to the microprocessor design 110, it should be appreciated that the techniques described herein may be applied more broadly to any kind of circuit design.

The techniques described above may be implemented, for example, in hardware, software, firmware, or any combination thereof. For example, the fitness calculator 202, configuration file manager 212, random code generator 112, simulator 114, and the configuration file synthesizer 222 may be implemented as computer programs. The techniques described above may be implemented in one or more computer programs executing on a programmable computer including a processor, a storage medium readable by the processor (including, for example, volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device. Program code may be applied to input entered using the input device to perform the functions described and to generate output. The output may be provided to one or more output devices.

Each computer program within the scope of the claims below may be implemented in any programming language, such as assembly language, machine language, a high-level procedural programming language, or an object-oriented programming language. The programming language may, for example, be a compiled or interpreted programming language.

Each such computer program may be implemented in a computer program product tangibly embodied in a machine-readable storage device for execution by a computer processor. Method steps of the invention may be performed by a computer processor executing a program tangibly embodied on a computer-readable medium to perform functions of the invention by operating on input and generating output. Suitable processors include, by way of example, both general and special purpose microprocessors. Generally, the processor receives instructions and data from a read-only memory and/or a random access memory. Storage devices suitable for tangibly embodying computer program instructions include, for example, all forms of non-volatile memory, such as semiconductor memory devices, including EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROMs. Any of the foregoing may be supplemented by, or incorporated in, specially-designed ASICs (application-specific integrated circuits). A computer can generally also receive programs and data from a storage medium such as an internal disk (not shown) or a removable disk. These elements will also be found in a conventional desktop or workstation computer as well as other computers suitable for executing computer programs implementing the methods described herein, which may be used in conjunction with any digital print engine or marking engine, display monitor, or other raster output device capable of producing color or gray scale pixels on paper, film, display screen, or other output medium.

CLAIMS

1. In a system (200) including a circuit design (110) and a plurality of circuit input vectors (102) which have been applied to the circuit design (110) to produce at least one occurrence of at least one event within the circuit design (110), a method comprising steps of:

- (A) identifying (326) at least one counter value indicating a number of occurrences of the at least one event; and
- (B) applying a fitness function (330) to the at least one counter value to produce an aggregate fitness value for the plurality of circuit input vectors.

2. The method of claim 1, wherein the at least one event comprises a plurality of events, wherein the step (A) comprises a step of:

(A) (1) identifying (324) a plurality of counter values indicating a number of occurrences of each of the plurality of events;

wherein the method further comprises a step of:

(C) prior to the step (B), initializing (322) the aggregate fitness value; and

wherein the step (B) comprises a step of performing the following steps for each event *e* in the plurality of events:

(B) (1) identifying (326) a counter value indicating a number of occurrences of event *e*;

(B) (2) applying (328) a counter fitness function to the counter value identified in step to produce a counter fitness value; and

(B) (3) applying (330) a combination function to the aggregate fitness value and the counter fitness value to produce a new value for the aggregate fitness value.

3. A system (200) comprising:
a circuit design (110);
a plurality of circuit input vectors (102) which have been applied to the circuit design (110) to produce at least one occurrence of at least one event within the circuit design (110);
counter value identification means (326) for identifying at least one counter value indicating a number of occurrences of the at least one event; and
fitness calculation means (330) for applying a fitness function to the at least one counter value to produce an aggregate fitness value for the plurality of circuit input vectors.

4. The system (2) of claim 3, wherein the at least one event comprises a plurality of events, wherein the counter value identification means comprises means for identifying (324) a plurality of counter values indicating a number of occurrences of each of the plurality of events, wherein the system further comprises means for initializing the aggregate fitness value, and wherein the fitness calculation means (330) comprises means for performing the following steps for each event *e* in the plurality of events: (1) identifying (326) a counter value indicating a number of occurrences of event *e*; (2) applying (328) a counter fitness function to the counter value identified in step to produce a counter fitness value; and (3) applying (330) a combination function to the aggregate fitness value and the counter fitness value to produce a new value for the aggregate fitness value.

5. In a system (200) including a circuit design (110) and a plurality of configuration files (108a-n), the plurality of configuration files including a plurality of knobs (104a-n) having a plurality of values specifying a plurality of probability distributions of a plurality of circuit input vectors, a method comprising steps of:

- (A) initializing (346) a synthesized configuration file;
- (B) selecting (342) a first subset of the plurality of configuration files as a potential parent pool; and
- (C) for each knob K_p (348) included in at least one of the plurality of configuration files in the potential parent pool, providing a synthesized value for knob K_p in the synthesized configuration file based on at least one value of knob K_p in the plurality of configuration files in the potential parent pool.

6. A system (200) comprising:
a circuit design (110);
a plurality of configuration files (108a-n), the plurality of configuration files including a plurality of knobs (104a-n) having a plurality of values specifying a plurality of probability distributions of a plurality of circuit input vectors;
means for initializing (346) a synthesized configuration file;
first selection means (342) for selecting a first subset of the plurality of configuration files as a potential parent pool; and
synthesizing means (348) for providing a synthesized value in the synthesized configuration file for each knob K_p included in at least one of the plurality of configuration files in the potential parent pool based on at least one value of knob K_p in the plurality of configuration files in the potential parent pool.

7. In a system (200) including a circuit design (110) and a plurality of configuration files (108a-n), the plurality of configuration files (108a-n) including a plurality of knobs (104a-n) having a plurality of values specifying a plurality of probability distributions of a plurality of circuit input vectors, a method comprising steps of:

- (A) initializing (346) a synthesized configuration file;
- (B) selecting (342) a subset of the plurality of configuration files having the highest fitness values as a potential parent pool;
- (C) randomly (344) selecting a subset of the potential parent pool as a plurality of direct parents;
- (D) for each knob K_p (348) included in at least one of the direct parents, providing a synthesized value for knob K_p in the synthesized configuration file based on at least one value of knob K_p in the plurality of direct parents.

8. A system (200) comprising:
a circuit design (110);
a plurality of configuration files (108a-n), the plurality of configuration files (108a-n) including a plurality of knobs (104a-n) having a plurality of values specifying a plurality of probability distributions of a plurality of circuit input vectors;
means for initializing (346) a synthesized configuration file;
means for selecting (342) a subset of the plurality of configuration files having the highest fitness values as a potential parent pool;
means for randomly selecting (344) a subset of the potential parent pool as a plurality of direct parents;
synthesizing means for providing (348), for each knob K_p included in at least one of the direct parents, a synthesized value for knob K_p in the synthesized configuration file based on at least one value of knob K_p in the plurality of direct parents.

9. In a system (200) including a circuit design (110) and a plurality of configuration files (108a-n), the plurality of configuration files (108a-n) including a plurality of knobs (104a-n) having a plurality of values specifying a plurality of probability distributions of a plurality of circuit input vectors, a method comprising steps of:

- (A) generating (302) a first synthesized configuration file based on the plurality of configuration files (108a-n), the first synthesized configuration file specifying a first synthesized probability distribution of a first subset of the plurality of circuit input vectors;
 - (B) using (304) the first synthesized configuration file to generate a first test case comprising the first subset of the plurality of circuit input vectors;
 - (C) providing (306) the first subset of the plurality of circuit input vectors as inputs to the circuit design;
 - (D) generating (308) a fitness value based on at least one event counter in the microprocessor design, the at least one event counter comprising a count of at least one occurrence of at least one event in the microprocessor design in response to provision of the first subset of the plurality of circuit input vectors;
 - (E) generating (310) a graded configuration file including the synthesized configuration file and the fitness value;
- and

(F) adding (312) the graded configuration file to the plurality of configuration files.

10. A system (200) comprising:

a circuit design (110);

a plurality of configuration files (108a-n), the plurality of configuration files (108a-n) including a plurality of knobs (104a-n) having a plurality of values specifying a plurality of probability distributions of a plurality of circuit input vectors;

means for generating (3020) a first synthesized configuration file based on the plurality of configuration files, the first synthesized configuration file specifying a first synthesized probability distribution of a first subset of the plurality of circuit input vectors;

means for using (304) the first synthesized configuration file to generate a first test case comprising the first subset of the plurality of circuit input vectors;

means for providing (306) the first subset of the plurality of circuit input vectors as inputs to the circuit design;

means for generating (308) a fitness value based on at least one event counter in the microprocessor design, the at least one event counter comprising a count of at least one occurrence of at least one event in the microprocessor design in response to provision of the first subset of the plurality of circuit input vectors;

means for generating (310) a graded configuration file including the synthesized configuration file and the fitness value; and

means for adding (312) the graded configuration file to the plurality of configuration files.



Application No: GB 0315869.8
Claims searched: 1 to 10

Examiner: Dean Parry
Date of search: 29 October 2003

Patents Act 1977 : Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
X	1, 3, and 5 to 10 at least	US 2002065640 A (RAMSEY) see whole document.
X	1, 3, and 5 to 10 at least	AN - 6622508 TI - Automatic test bench generation for validation of RT-level descriptions: an industrial experience. PUB - 2000, Los Alamitos, CA, USA, IEEE Comput. Soc, USA. CONF - Proceedings of Meeting on Design Automation and Test in Europe, Paris, France, 27-30 March 2000. AU - Como F ; Sonza Reorda M ; Squillero G ; Manzone A ; Pincetti A. ISBN 0-7695-0537-6 ; PG - 385 - 389 ; DT - CA Conference Article ; ED - Marwedel P ; Bolsens I ; ORD - 2000-00-00.

Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKCV:

G4A

Worldwide search of patent documents classified in the following areas of the IPC⁷:

G06F

The following online and other databases have been used in the preparation of this search report:

Online : EPODOC, JAPIO, WPI, FULL TEXT, INSPEC